Dr. Marcus L. Graham, Dr. Roy Streit, Michael Walsh, NUWCDIVNPT

A Strategy for Introducing Automation and Gaining Operator Confidence – A Case Study

ABSTRACT

The goal of automation in submarine sonar is to alleviate operator workload to enable the sonar operator to focus on other tasks. In the development of detection systems, engineers are universally trained to "optimize" the system set point by choosing a threshold on a curve that characterizes the trade-off between the probability of detecting a desired event and the probability of false detections. Such thresholds are widely accepted and often used to highlight "low level" sensor data that the operator expects to treat manually.

However, for sufficiently important or "higher level" tasks, the occasional presence of false events in the stream of automated detections often requires the operator to validate to *all* reported detections. While this may reduce the operator workload, it may also diminish operator confidence in the automation. Erosion of confidence can quickly lead to the rejection of a useful product, especially in the early introductory phase.

An alternative strategy for introducing automation, especially suitable for high level tasking, is to choose a set point that virtually eliminates false detections at the expense of missed true event detections. This strategy may facilitate technology acceptance by the operator. This seems to be due to the fact that operators develop confidence in the automation tool because they learn that they never have to validate any of the automated detections -- because they are nearly always correct. The result is a true reduction of operator workload, even though the operator still has workload remaining.

These lessons were learned while introducing an automation tool for re-assigning acoustic contact trackers after data gaps produced by ownship turns. Details of this example will be described, and the solution that resulted in general acceptance of the automation tool will be outlined. The concept of operation for the resulting automation product will be given and representations of automation certainty discussed.

INTRODUCTION

Event detection is a major component in many information-processing problems. Difficult detection problems requiring evaluation of complex criteria or aggregation of evidence from multiple sources often require operator interactive methods. With ever increasing data volume due to improved sensors and increasing processing rates, it is increasingly common for operators to become overloaded.

In a system that relies on an operator for detection decisions the workload can often be triaged into two kinds of decisions – those that are relatively easy for the operator to assess but still require operator action, and those that are difficult to assess and require more deliberation. It is obvious that decisions of the second kind (those requiring deliberation) pose a burden on operators. Both kinds, however, can be significant components of the total operator workload.

Decisions of the first kind (those requiring little or no deliberation) can pose a significant operator burden because there may be a large number of them. The decision criteria in these cases may often be easily assessed and may have decision statistics well above (or below) some statistical decision threshold. Nonetheless, the operator must still cycle through one or more menus to actuate (implement) these decisions in the system. From a systems engineering perspective this is a poor utilization of a precious resource – the operator. The operator's time is better devoted to the more challenging cases that require training, experience, and reasoning.

It is not the perspective of this paper that all high-level tasks should be automated. For very high-level tasks, where decisions are time critical, or where the cost of false alarms and missed detections may be catastrophic, it may be necessary for operator intensive modes of operation to persist. Attention here focuses on only those tasks for which the community deems automation appropriate.

AUTOMATION INSERTION

From a traditional statistical perspective it is natural to analyze a detection problem in terms of the trade-off between probability of detection and the probability of false alarm (PD/PFA). The common tendency is to view the problem as a completely binary decision, with the only options being detect or no-detect. It is important to realize that a third state, that of ignorance or indecision, is viable and in fact can sometimes be more useful in practice than a "confidence factor" attached to an automated detect or no-detect decision.

The addition of a third, in-between state to the automation string is significant in practice. False detections and missed detections often correspond to numerical values of the decision statistic near the decision threshold. By adding an third decision state, these harder, less certain cases can be flagged for further analysis by the operator, while the easier, more certain detections and rejections can be automated. The automation alerts the operator only to those decisions in which it finds itself in this third state of indecision.

When introducing automation in a high-level critical operation that is currently a manually intensive task, and where operators must maintain some degree of control, it is critically important to build operator confidence. From this standpoint it is advantageous to consider different utility functions for the automation output at different stages of development.

When automation technology is first introduced, relatively few false alarms (or missed detections, depending on the problem) may have an unexpected and far-reaching impact. Such errors, even though they may be infrequent, may nonetheless be sufficient to kill a promising detection technology. This is because each false alarm/missed detection increases operator doubt regarding *every* subsequent automated detect/no-detect decision. As a result the operator may quickly perceive a need to validate all automated decisions and, hence, there will be no reduction in workload in practice. During technology insertion the utility function must account for the operator's skepticism and place a much higher cost on false alarms and missed detections than on a reported state of indecision.

As automation technology improves and matures, it may eventually perform as well as the operator. If such should turn out to be the case, the proposed third state may ultimately prove to be of only nominal value. Reporting this third state would seem to be advantageous primarily during the initial automation technology insertion phase, when automation performance is being closely monitored in-situ by operators.

The need for a third, uncertain state in the decision space was highlighted by recent experience with trying to introduce an automation tool into the current processing string of ASTO's Advanced Processing Build (APB). The example is outlined in the next section. This three-state decision space is somewhat more complicated than the traditional two-state model of detection/no-detection, but the basic points illustrate significant value added.

AN ILLUSTRATIVE EXAMPLE

These lessons were learned while attempting to introduce an automation tool for re-associating towed array acoustic contact trackers after data gaps produced by ownship turns. The problem arises due to the observability limitations of angle-only data from a towed line array. Most often, the contact to be tracked is assumed to be traveling at a constant velocity during the observation interval.

The line array yields bearing-like data that lie on cones whose axis is the array axis as shown in figure 1 (Graham et al. 1994). Sequential observations of this cone angle during a constant velocity tow yield sufficient information to estimate the range-normalized parameters that characterize the contact's location and motion, but not the contact range itself (Hammel and Aidala, 1985). When the observing platform maneuvers, the target range becomes theoretically observable from the cone angle sequence.

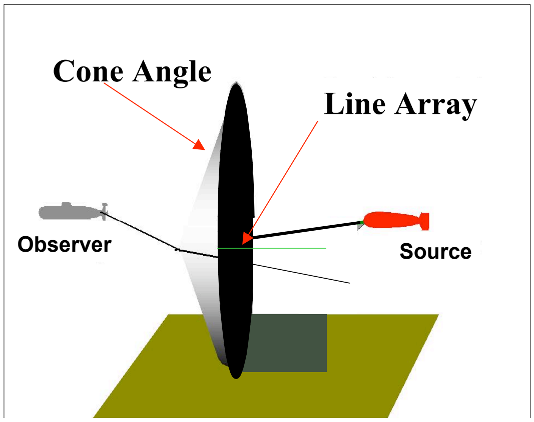

The difficulty is that during the maneuver the array becomes severely distorted, and it is temporarily unable to produce reliable cone-angle or other geometrically useful data. Further, due to the ambiguity of the cone angle measurement (the circumferential angle is unobservable), it is impossible to predict where to look for the corresponding data after the maneuver when the array finally stabilizes without knowledge of the contact range.

Thus, there is a need to associate cone angle sequences from before and after the ownship maneuver in order to estimate the contact range. This association was traditionally done by the operator, and it is complicated by the presence of multiple contacts in the surrounding waters. The sonar operators typically rely on ancillary information to determine which bearing sequence from before the maneuver goes with each bearing sequence after the maneuver. A key piece of association evidence is found in the frequency content of the received energy. Sonar operators comb through large displays of the acoustic energy on detected cone angle sequences to search for matching frequency characteristics. When large numbers of targets are present this can be and is a daunting task.

Recent introduction of spectrally based tracking methods have set the stage for automating the re-association task. These methods provide a nonparametric estimate of both the target and noise spectra, and facilitate the formation of a target-specific detection statistic. This statistic, termed the cross-ratio (XR), provides a simple and robust approximation to the optimal detection statistic (Graham, Walsh and Streit, 2003) for the matching target spectra.

The spectral estimates from data gathered prior to the observer maneuver are used in formulating the XR statistic to be applied after the maneuver. Peak values of the XR statistic are examined for consistency across time to arrive at a Best Beam Indicator (BBI) for locating the new cone angle of the target. During periods when the statistic is not consistent across time, no BBI is provided, indicating uncertainty in the target location in the sensor beams.

Initial testing indicated excellent performance, providing the correct pre- and post-maneuver associations for most contacts present. Of course, it is not 100% accurate all the time. Nonetheless, the tool alleviates the bulk of the operator's workload, allowing them to focus on the more difficult cases where the XR is inconsistent and no BBI provided.

ownship maneuver, and he wants to know which data is best associated with it after the maneuver. (During the maneuver, as may be seen, the data is so poorly defined that trackers lose track and fly unnaturally across many beams).

The proposed BBI shows as a cursor at the top of the waterfall data plot pointing the operator where to restart his selected track. At the operator's discretion, he is also able to bring up a plot of the XR statistic itself for the same time period to aid him in cases of difficult assessment. The BBI Algorithm

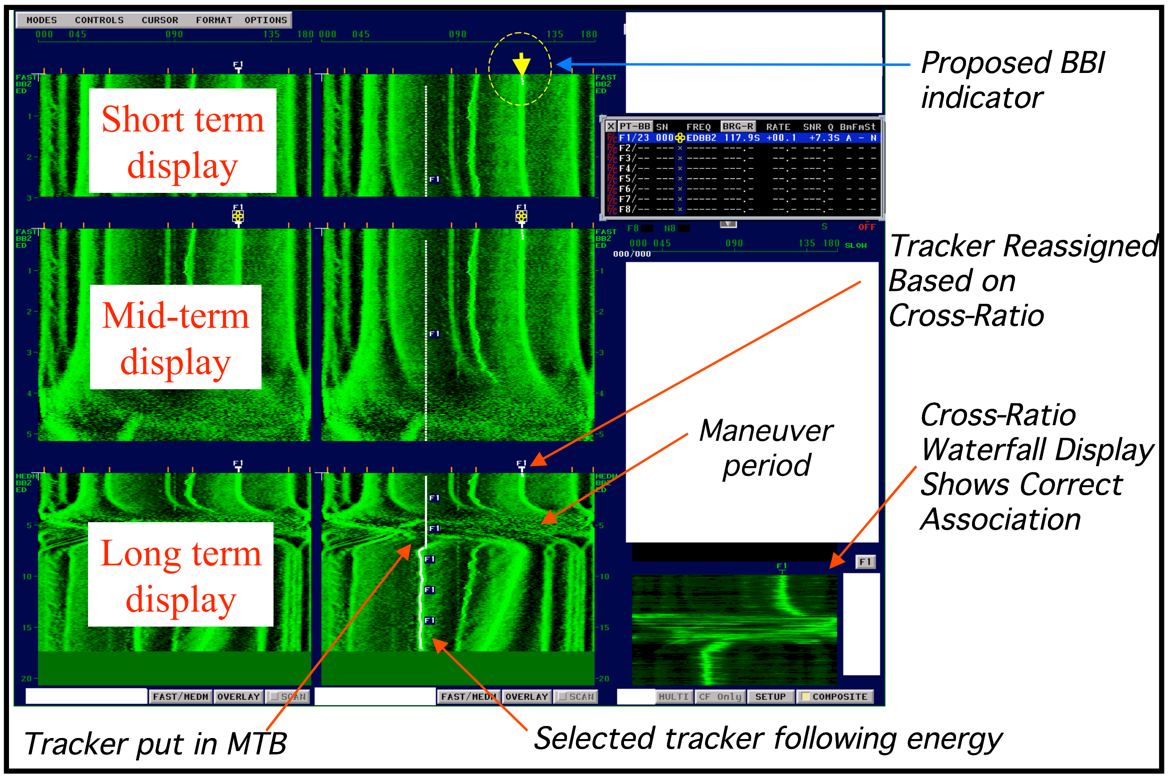

Figure 2. Prototype Display

A prototype display is shown in figure 2, where the plot is of time versus cone-angle beam and the intensity of the green color displayed indicates the level of energy on that beam at that time. The plot window is of fixed duration, and the most recent time is at the top – thus the data "waterfalls" down the screen as new data scans arrive. The operator selects a track existing prior to the

was proposed for inclusion in the ASTO sponsored Advanced Processor Build - (Acoustic) (APB(A)) and thus was required to undergo independent testing.

Independent testing was conducted at MIT/LL by individuals experienced in the introduction of sonar detection aids (also sometimes known as Bell-Ringers). In the first phase of testing it was noted that the occasional miss-cue by the algorithm degrade its acceptance due to the operators need to check all cases. However, with a simple pre-screening function applied to eliminate problematic cases (cases where multiple contact spectra are not sufficiently distinct prior to the maneuver), an algorithm with 100% correct identification on those cases passing the screen could be achieved (Rholt, Private Communication 2003).

It was determined that for the cases tested, approximately 60% passed the screen and of those 100% were correctly associated across pre and post maneuver periods. (Payne, W.H.) While reporting on only 60% of the cases may seem unusually low, the resulting 100% accuracy in testing resulted in significantly more favorable acceptance by those with operator experience. Reporting on a higher percentage of cases, but with anything less than perfect decisions was not viewed as beneficial to the operator due to the perceived need to check all cases.

CONCLUSION

When introducing technology for high-level decisions, the cost of false alarms must be carefully weighted. The idea that an optimal operating point may exist with anything less than 100% accuracy must be contrasted with the need to gain operator acceptance of the technology. In cases where operators have historically provided the functionality, it may be better to chose a conservative threshold and allow the operators to continue to fulfill the function for difficult cases. This builds operator trust and permits continued technology development, whereupon operating points can be adjusted as automation technology matures. Lack of operator trust may result in premature technology rejection and thus kill development cycles, yielding either an inferior overall system or significantly delaying the introduction of desired automation technology.

REFERENCES

Graham, M.L., Santos, M.C. Hammel, S.E. and Pendergrass, N.A., Two Compensation Methods for Sensor Orientation Uncertainty in Target Motion Analysis, NUWC-NPT Technical Report 10,402, 15 September 1994.

Hammel, S.E. and Aidala, V.J., "Observability Requirements for Three Dimensional Tracking via Angle Measurements," IEEE Tran. On Aerospace and Electronic Systems, Vol AES-21, No. 2, March 1985.

Graham, M.L., Walsh, M.J., and Streit, R.L., Optimum Nonlinear Filter for Target Detection and Interference Rejection Proceedings of Sensors & Electronics Technology Panel Symposium:Target Tracking And Data Sensor Fusion For Military Observation Systems, Budapest, Hungary, 15-17 October 2003.

Rholt, K., MIT/LL, Private communication, August 2003.

Payne, W.H., Presentation to the Advanced Processor Build Acoustic Automation Working Group (APB(A)-AWG), 27 January, 2004.

ACKNOWLEDGMENTS

This work was supported by J. Tague, ONR 321, under the Sonar Automation Technology Task

---

Dr. Marcus L. Graham is an Electrical Engineer at NUWCDIVNPT, Code 2211, (401)-832-3846, E-mail: GrahamML@npt.nuwc.navy.mil

Dr. Roy L. Streit is a Mathematician at NUWCDIVNPT, Code 2002, (401)-832-4189, E-mail: StreitRL@npt.nuwc.navy.mil

Michael J. Walsh is an Electrical Engineer at NUWCDIVNPT, Code 2211, (401)-832-4155, E-mail: WalshMJ@npt.nuwc.navy.mil